Package: gipsDA (via r-universe)

November 1, 2026

Title Discriminant Analysis with Permutation-Invariant Covariance Models

Version 1.0.0

URL https://antonikingston.github.io/gipsDA/, https://github.com/AntoniKingston/gipsDA

BugReports https://github.com/AntoniKingston/gipsDA/issues

Description Extends classical linear and quadratic discriminant analysis by incorporating permutation-group symmetries into covariance matrix estimation. Methods based on the 'gips' framework identify and impose permutation structures that regularize covariance estimates and improve stability and interpretability for symmetric or exchangeable features. The package provides pooled and class-specific covariance models, including multi-class variants with shared or independently estimated symmetry structures. The underlying methodology is described by Graczyk et al. (2022) <doi:10.1214/22-AOS2174> and Chojecki, Morgen, and Kołodziejek (2025) <doi:10.18637/jss.v112.i07>.

License GPL (>= 3)

Imports gips (>= 1.3.0), jsonlite, lattice, MASS, stats, stringi

Suggests knitr, mockery, rmarkdown, testthat (>= 3.0.0), roxygen2, withr

Depends R (>= 4.1.0)

Encoding UTF-8

Roxygen list(markdown = TRUE)

RoxygenNote 8.0.0

VignetteBuilder knitr

Config/testthat/edition 3

Config/Needs/website pkgdown

Config/pak/sysreqs libgmp3-dev libicu-dev

Repository https://antonikingston.r-universe.dev

Date/Publication 2026-10-01 22:15:21 UTC

RemoteUrl https://github.com/antonikingston/gipsda

RemoteRef HEAD

RemoteSha 1a54cbb2675e475d55ef3d7aeb0d75fc08187fac

Contents

coef.gipslda 2
gipslda 3
gipsmultqda 5
gipsqda 8
pairs.gipslda 10
plot.gipslda 11
predict.gipslda 12
predict.gipsmultqda 13
predict.gipsqda 14
print.gipslda 15
summary.gipslda 15

Index **17**

`coef.gipslda` *Extract gips LDA discriminant coefficients*

Description

Extract gips LDA discriminant coefficients

Usage

```
## S3 method for class 'gipslda'
coef(object, ...)
```

Arguments

`object`	A fitted "`gipslda`" object.
`...`	Unused.

Value

The linear discriminant transformation matrix.

See Also

`gipslda`

gipslda *Linear Discriminant Analysis with gips Covariance Projection*

Description

Linear discriminant analysis (LDA) using covariance matrices projected via the *gips* framework to enforce permutation symmetry and improve numerical stability.

Usage

```
gipslda(x, ...)

## S3 method for class 'formula'
gipslda(formula, data, ..., subset, na.action)

## Default S3 method:
gipslda(x, grouping, prior = proportions,
  tol = 1e-4, weighted_avg = FALSE,
  MAP = TRUE, optimizer = NULL, max_iter = NULL,
  show_progress_bar = FALSE, store_probabilities = TRUE, ...)

## S3 method for class 'data.frame'
gipslda(x, ...)

## S3 method for class 'matrix'
gipslda(x, grouping, ..., subset, na.action)
```

Arguments

`x`	(required if no formula is given as the principal argument) a matrix or data frame or Matrix containing the explanatory variables.
`...`	Arguments passed to or from other methods.
`formula`	A formula of the form `groups ~ x1 + x2 + ...`. The response is the grouping factor and the right-hand side specifies the (non-factor) discriminators.
`data`	An optional data frame, list or environment from which variables specified in `formula` are preferentially taken.
`grouping`	(required if no formula principal argument is given) a factor specifying the class for each observation.
`prior`	The prior probabilities of class membership. If unspecified, the class proportions for the training set are used.
`tol`	A tolerance to decide if a matrix is singular; variables whose variance is less than `tol^2` are rejected.
`subset`	An index vector specifying the cases to be used in the training sample. (NOTE: must be named.)
`na.action`	A function specifying the action for `NA`s.

MAP	Logical; whether to compute a Maximum A Posteriori gips projection. If FALSE, projected matrices are averaged using retained posterior permutation probabilities.
optimizer	Character; optimization method used by gips (e.g. "BF" or "MH").
max_iter	Maximum number of iterations for the optimizer.
show_progress_bar	Logical; if TRUE, display the progress bar from the underlying gips optimizer. Defaults to FALSE.
store_probabilities	Logical; if TRUE, store estimated posterior probabilities of retained permutations. If FALSE and MAP = TRUE, only the selected MAP permutation is stored. Defaults to TRUE.
weighted_avg	Logical; if FALSE, use the pooled within-class scatter matrix. If TRUE, use a class-proportion-weighted average of class-specific covariance matrices.

Details

This function is a minor modification of lda, replacing the classical sample covariance estimators by projected covariance matrices obtained using the gips framework.

Unlike classical LDA, the within-class covariance matrix is first projected onto a permutation-invariant structure using the gips framework. This can stabilize covariance estimation in high dimensions or when symmetry assumptions are justified.

The choice of optimizer and MAP estimation affects both the covariance estimate and the resulting discriminant directions.

See Chojecki et al. (2025) for theoretical background.

Value

An object of class "gipslda" containing:

- prior: prior class probabilities
- counts: number of observations per class
- means: group means
- scaling: transformation matrix of linear discriminants
- lev: class labels
- svd: singular values of the between-class scatter
- N: number of observations
- optimization_info: estimated posterior probabilities of retained permutations if stored; otherwise NULL
- selected_map_permutation: MAP permutation selected by the gips optimization and used for MAP covariance projection
- call: matched call
- Formula fits additionally contain terms, contrasts, xlevels, and any recorded na.action.

Note

This function is inspired by lda but is not a drop-in replacement. The covariance estimator, optimization procedure, and returned object differ substantially.

References

Chojecki, A., et al. (2025). *Learning Permutation Symmetry of a Gaussian Vector with gips in R*. Journal of Statistical Software, **112**(7), 1–38. doi:10.18637/jss.v112.i07

See Also

lda, gips

Examples

```
Iris <- data.frame(rbind(iris3[, , 1], iris3[, , 2], iris3[, , 3]),
  Sp = rep(c("s", "c", "v"), rep(50, 3))
)
train <- sample(1:150, 75)
z <- gipslda(Sp ~ ., Iris, prior = c(1, 1, 1) / 3, subset = train)
predict(z, Iris[-train, ])$class
(z1 <- update(z, . ~ . - Petal.W.))
```

---

gipsmultqda *Quadratic Discriminant Analysis with multiple gips-projected covariances*

Description

Quadratic Discriminant Analysis (QDA) in which each class covariance matrix participates in one joint *gips* projection, allowing the covariance matrices to share an estimated permutation symmetry.

Usage

```
gipsmultqda(x, ...)

## S3 method for class 'formula'
gipsmultqda(formula, data, ..., subset, na.action)

## Default S3 method:
gipsmultqda(x, grouping, prior = proportions,
  nu = 5, MAP = TRUE, optimizer = NULL, max_iter = NULL,
  show_progress_bar = FALSE, store_probabilities = TRUE, ...)

## S3 method for class 'data.frame'
gipsmultqda(x, ...)
```

```
## S3 method for class 'matrix'
gipsmultqda(x, grouping, ..., subset, na.action)
```

Arguments

x	(required if no formula is given as the principal argument) a matrix or data frame containing the explanatory variables.
...	Arguments passed to or from other methods.
formula	A formula of the form groups ~ x1 + x2 + ... The response is the grouping factor and the right-hand side specifies the (non-factor) discriminators.
data	An optional data frame, list or environment from which variables specified in formula are preferentially taken.
grouping	(required if no formula is given) a factor specifying the class for each observation.
prior	Prior probabilities of class membership. If omitted, training-set class proportions are used. Supplied values must sum to one.
nu	Reserved for compatibility with related QDA interfaces; it is not currently used by the covariance projection.
MAP	Logical; if TRUE, a maximum a posteriori covariance projection is used. If FALSE, projections are averaged using retained posterior permutation probabilities.
optimizer	Character string specifying the optimization method used for covariance projection. If NULL, a default choice is made based on the problem dimension.
max_iter	Maximum number of iterations for stochastic optimizers.
show_progress_bar	Logical; if TRUE, display the progress bar from the underlying gips optimizer. Defaults to FALSE.
store_probabilities	Logical; if TRUE, store estimated posterior probabilities of retained permutations. If FALSE and MAP = TRUE, only the selected MAP permutation is stored. Defaults to TRUE.
subset	An index vector specifying the cases to be used in the training sample. (NOTE: must be named.)
na.action	A function specifying the action to be taken if NAs are found.

Details

This function is a modification of qda in which the class-specific covariance matrices are jointly projected to improve numerical stability and exploit shared symmetry assumptions.

In contrast to classical QDA, which estimates each class covariance matrix independently, gipsmultqda performs a joint projection of all class covariance matrices using gips. This allows the incorporation of shared permutation symmetries and can improve classification performance in high-dimensional or small-sample regimes.

Several classification rules are available via predict.gipsmultqda, including plug-in, predictive, debiased, and leave-one-out cross-validation.

Value

An object of class "gipsmultqda" containing:

- prior: prior probabilities of the groups
- counts: number of observations per group
- means: group means
- scaling: array of group-specific scaling matrices derived from the projected covariance matrices
- ldet: log-determinants of the projected covariance matrices
- lev: class labels
- N: total number of observations
- optimization_info: estimated posterior probabilities of retained permutations from the joint gips optimization if stored; otherwise NULL
- selected_map_permutation: MAP permutation selected by the joint gips optimization and used for MAP covariance projection
- call: the matched call
- Formula fits additionally contain terms, contrasts, xlevels, and any recorded na.action.

Note

This function is not a drop-in replacement for qda. The covariance estimation, returned object, and classification rules differ substantially.

The theoretical background and details of covariance projection are documented by gips and Chojecki et al. (2025).

See Also

qda, predict.gipsmultqda, gipsqda, gipslda

Examples

```
tr <- sample(1:50, 25)
train <- rbind(iris3[tr, , 1], iris3[tr, , 2], iris3[tr, , 3])
test <- rbind(iris3[-tr, , 1], iris3[-tr, , 2], iris3[-tr, , 3])
cl <- factor(c(rep("s", 25), rep("c", 25), rep("v", 25)))
z <- gipsmultqda(train, cl)
predict(z, test)$class
```

gipsqda *Quadratic Discriminant Analysis with gips covariance projection*

Description

Quadratic discriminant analysis (QDA) using covariance matrices projected via the *gips* framework to enforce permutation symmetry and improve numerical stability.

Usage

```
gipsqda(x, ...)

## S3 method for class 'formula'
gipsqda(formula, data, ..., subset, na.action)

## Default S3 method:
gipsqda(x, grouping, prior = proportions,
  nu = 5, MAP = TRUE, optimizer = NULL, max_iter = NULL,
  show_progress_bar = FALSE, store_probabilities = TRUE, ...)

## S3 method for class 'data.frame'
gipsqda(x, ...)

## S3 method for class 'matrix'
gipsqda(x, grouping, ..., subset, na.action)
```

Arguments

x	(required if no formula is given as the principal argument) a matrix or data frame containing the explanatory variables.
...	Arguments passed to or from other methods.
formula	A formula of the form groups ~ x1 + x2 + ... The response is the grouping factor and the right-hand side specifies the (non-factor) discriminators.
data	An optional data frame, list or environment from which variables specified in formula are preferentially taken.
grouping	(required if no formula is given) a factor specifying the class for each observation.
prior	The prior probabilities of class membership. If omitted, training-set class proportions are used. Supplied values must sum to one and match the number of groups.
nu	Reserved for compatibility with related QDA interfaces; it is not currently used by the covariance projection.
MAP	Logical; if TRUE, maximum a posteriori covariance projection is used. If FALSE, projections are averaged using retained posterior permutation probabilities.

optimizer	Character string specifying the optimization method used for covariance projection. If NULL, a default choice depending on the problem dimension is used.
max_iter	Maximum number of iterations for stochastic optimizers.
show_progress_bar	Logical; if TRUE, display the progress bar from the underlying gips optimizer. Defaults to FALSE.
store_probabilities	Logical; if TRUE, store estimated posterior probabilities of retained permutations. If FALSE and MAP = TRUE, only the selected MAP permutations are stored. Defaults to TRUE.
subset	An index vector specifying the cases to be used in the training sample. (NOTE: must be named.)
na.action	A function specifying the action to be taken if NAs are found.

Details

This function is a minor modification of qda, replacing the classical sample covariance estimators by projected covariance matrices obtained using the gips framework.

Quadratic discriminant analysis models each class with its own covariance matrix. In gipsqda, these covariance matrices are projected using the *gips* framework, which enforces permutation symmetry and mitigates singularity and overfitting in high-dimensional or small-sample settings.

Classification can be performed using plug-in, predictive, debiased, or leave-one-out cross-validation rules via predict.gipsqda.

Value

An object of class "gipsqda" containing the following components:

- prior: prior probabilities of the groups
- counts: number of observations in each group
- means: group means
- scaling: group-specific scaling matrices derived from the projected covariance matrices
- ldet: log-determinants of the projected covariance matrices
- lev: class labels
- N: total number of observations
- optimization_info: named list of estimated posterior probabilities of retained permutations, one element per class, if stored; otherwise a named list of NULL values
- selected_map_permutation: named list of MAP permutations selected independently for each class
- call: the matched call
- Formula fits additionally contain terms, contrasts, xlevels, and any recorded na.action.

Note

The function may be called with either a formula interface or with a matrix and grouping factor. Arguments subset and na.action, if used, must be named.

References

Chojecki, A., et al. (2025). *Learning Permutation Symmetry of a Gaussian Vector with gips in R*. Journal of Statistical Software, **112**(7), 1–38. doi:10.18637/jss.v112.i07

Venables, W. N. and Ripley, B. D. (2002). *Modern Applied Statistics with S*. Fourth edition. Springer.

See Also

qda, predict.gipsqda, gipslda, lda

Examples

```
tr <- sample(1:50, 25)
train <- rbind(iris3[tr, , 1], iris3[tr, , 2], iris3[tr, , 3])
test <- rbind(iris3[-tr, , 1], iris3[-tr, , 2], iris3[-tr, , 3])
cl <- factor(c(rep("s", 25), rep("c", 25), rep("v", 25)))
z <- gipsqda(train, cl)
predict(z, test)$class
```

---

pairs.gipslda *Pairs plot for a fitted gips LDA model*

Description

Displays discriminant coordinates using either base graphics or a lattice scatterplot matrix.

Usage

```
## S3 method for class 'gipslda'
pairs(
  x,
  labels = colnames(x),
  panel = panel.gipslda,
  dimen,
  abbrev = FALSE,
  ...,
  cex = 0.7,
  type = c("std", "trellis")
)
```

Arguments

x	A fitted "gipslda" object.
labels	Labels for discriminant dimensions.
panel	Panel function used by base graphics.

dimen	Maximum number of discriminant dimensions to display.
abbrev	Logical or integer controlling abbreviation of class labels.
...	Additional graphical arguments.
cex	Character expansion used by the default panel.
type	Either "std" for base graphics or "trellis" for a lattice display.

Value

Invisibly returns NULL.

See Also

gipslda, plot.gipslda

---

plot.gipslda *Plot a fitted gips LDA model*

Description

Displays training observations in discriminant space. One discriminant dimension is shown as class-wise histograms or densities, two dimensions as an equal-scale scatter plot, and higher dimensions as a pairs plot.

Usage

```
## S3 method for class 'gipslda'
plot(
  x,
  panel = panel.gipslda,
  ...,
  cex = 0.7,
  dimen,
  abbrev = FALSE,
  xlab = "LD1",
  ylab = "LD2"
)
```

Arguments

x	A fitted "gipslda" object.
panel	Panel function used for scatter plots.
...	Additional graphical arguments.
cex	Character expansion used by the default panel.
dimen	Maximum number of discriminant dimensions to display.
abbrev	Logical or integer controlling abbreviation of class labels.
xlab, ylab	Axis labels.

Value

Invisibly returns NULL.

See Also

gipslda, pairs.gipslda

---

predict.gipslda *Predict from a gips LDA model*

Description

Computes class assignments, posterior probabilities, and linear discriminant coordinates for new observations.

Usage

```
## S3 method for class 'gipslda'
predict(
  object,
  newdata,
  prior = object$prior,
  dimen,
  method = c("plug-in", "predictive", "debiased"),
  ...
)
```

Arguments

object	A fitted "gipslda" object.
newdata	An optional matrix or data frame of observations. If omitted, the training data are reconstructed from the fitted call.
prior	Prior class probabilities used for prediction.
dimen	Number of discriminant dimensions to use. The default uses all dimensions available in object.
method	Prediction rule: "plug-in", "predictive", or "debiased".
...	Further arguments passed to or from methods.

Value

A list with class (predicted classes), posterior (posterior class probabilities), and x (linear discriminant coordinates).

See Also

gipslda, predict.lda

`predict.gipsmultqda` *Predict from a joint-projection gips QDA model*

Description

Computes class assignments and posterior probabilities from a fitted `"gipsmultqda"` model. Prediction rules match those of `predict.gipsqda`; the fitted covariance estimates differ because they were projected jointly.

Usage

```
## S3 method for class 'gipsmultqda'
predict(
  object,
  newdata,
  prior = object$prior,
  method = c("plug-in", "predictive", "debiased", "looCV"),
  ...
)
```

Arguments

`object`	A fitted `"gipsmultqda"` object.
`newdata`	An optional matrix or data frame of observations. Omit this argument when using `method = "looCV"`.
`prior`	Prior class probabilities used for prediction.
`method`	Prediction rule: `"plug-in"`, `"predictive"`, `"debiased"`, or `"looCV"`. Leave-one-out prediction classifies the reconstructed training observations.
`...`	Further arguments passed to or from methods.

Value

A list with `class` (predicted classes) and `posterior` (posterior class probabilities).

See Also

`gipsmultqda`, `predict.gipsqda`

`predict.gipsqda` *Predict from a gips QDA model*

Description

Computes class assignments and posterior probabilities from a fitted `"gipsqda"` model.

Usage

```
## S3 method for class 'gipsqda'
predict(
  object,
  newdata,
  prior = object$prior,
  method = c("plug-in", "predictive", "debiased", "looCV"),
  ...
)
```

Arguments

`object`	A fitted `"gipsqda"` object.
`newdata`	An optional matrix or data frame of observations. Omit this argument when using `method = "looCV"`.
`prior`	Prior class probabilities used for prediction.
`method`	Prediction rule: `"plug-in"`, `"predictive"`, `"debiased"`, or `"looCV"`. Leave-one-out prediction classifies the reconstructed training observations.
`...`	Further arguments passed to or from methods.

Value

A list with `class` (predicted classes) and `posterior` (posterior class probabilities).

See Also

`gipsqda`, `predict.qda`

print.gipslda *Print a fitted gipsDA model*

Description

Prints the main components of a fitted `gipsDA` model, including the model call, fitting options, group means, class counts, selected MAP permutation, and posterior probabilities of retained permutations when stored.

Usage

```
## S3 method for class 'gipslda'
print(x, ...)

## S3 method for class 'gipsqda'
print(x, ...)

## S3 method for class 'gipsmultqda'
print(x, ...)
```

Arguments

`x`	A fitted gipsDA model.
`...`	Further arguments passed to printing methods.

Value

Invisibly returns `x`.

Examples

```
fit <- gipslda(Species ~ ., data = iris, optimizer = "BF")
print(fit)
```

summary.gipslda *Summarize a fitted gipsDA model*

Description

Creates a compact summary object for a fitted `gipsDA` model. The summary contains the most important fitted quantities and can be printed using the corresponding `print()` method.

Usage

```
## S3 method for class 'gipslda'
summary(object, ...)

## S3 method for class 'gipsqda'
summary(object, ...)

## S3 method for class 'gipsmultqda'
summary(object, ...)
```

Arguments

object	A fitted gipsDA model.
...	Further arguments passed to or from methods.

Value

An object of class "summary.gipsda". Common components include:

- model: model type.
- call: original model call.
- n: number of observations used for fitting.
- p: number of predictors.
- groups: fitted class labels.
- counts: class counts.
- prior: prior class probabilities.
- means: group means.
- fit_info: fitting options, including MAP, optimizer, max_iter, store_probabilities, and for LDA also weighted_avg.
- optimization_info: stored posterior probabilities of retained permutations, if available.
- selected_map_permutation: selected MAP permutation or permutations.

LDA summaries additionally contain scaling, svd, and proportion_trace. QDA summaries additionally contain scaling_dim and ldet.

Examples

```
fit <- gipslda(Species ~ ., data = iris, optimizer = "BF")
summary(fit)

summary_object <- summary(fit)
names(summary_object)
```

Index

∗ **classification**
gipsmultqda, 5
gipsqda, 8
∗ **multivariate**
gipslda, 3
gipsmultqda, 5
gipsqda, 8

coef.gipslda, 2

gips, *4–7*, *9*
gipslda, *2*, 3, *7*, *10–12*
gipsmultqda, 5, *13*
gipsqda, *7*, 8, *14*

lda, *4, 5*, *10*

model.frame.gipslda *(*gipslda*)*, 3
model.frame.gipsmultqda *(*gipsmultqda*)*, 5
model.frame.gipsqda *(*gipsqda*)*, 8

pairs.gipslda, 10, *12*
plot.gipslda, 11, *11*
predict.gipslda, 12
predict.gipsmultqda, *6, 7*, 13
predict.gipsqda, *9, 10*, *13*, 14
predict.lda, *12*
predict.qda, *14*
print.gipslda, 15
print.gipsmultqda *(*print.gipslda*)*, 15
print.gipsqda *(*print.gipslda*)*, 15

qda, *6, 7*, *9, 10*

summary.gipslda, 15
summary.gipsmultqda *(*summary.gipslda*)*, 15
summary.gipsqda *(*summary.gipslda*)*, 15